

# Improving Tor using a TCP-over-DTLS Tunnel

Joel Reardon\*  
Google Switzerland GmbH  
Brandschenkestrasse 110  
Zürich, Switzerland  
reardon@google.com

Ian Goldberg  
University of Waterloo  
200 University Ave W.  
Waterloo, ON, Canada  
iang@cs.uwaterloo.ca

## Abstract

The Tor network gives anonymity to Internet users by relaying their traffic through the world over a variety of routers. All traffic between any pair of routers, even if they represent circuits for different clients, are multiplexed over a single TCP connection. This results in interference across circuits during congestion control, packet dropping and packet reordering. This interference greatly contributes to Tor’s notorious latency problems.

Our solution is to use a TCP-over-DTLS (Datagram Transport Layer Security) transport between routers. We give each stream of data its own TCP connection, and protect the TCP headers—which would otherwise give stream identification information to an attacker—with DTLS. We perform experiments on our implemented version to illustrate that our proposal has indeed resolved the cross-circuit interference.

## 1 Introduction

*Tor* [3] is a tool to enable Internet privacy that has seen widespread use and popularity throughout the world. Tor consists of a network of thousands of nodes—known as Onion Routers (ORs)—whose operators have volunteered to relay Internet traffic around the world. Clients—known as Onion Proxies (OPs)—build circuits through ORs in the network to dispatch their traffic. Tor’s goal is to frustrate an attacker who aims to match up the identities of the clients with the actions they are performing. Despite its popularity, Tor has a problem that dissuades its ubiquitous application—it imposes greater latency on its users than they would experience without Tor.

While some increased latency is inevitable due to the increased network path length, our experiments show that this effect is not sufficient to explain the increased cost. In Section 2 we look deeper, and find a component of the transport layer that can be changed to improve Tor’s performance. Specifically, each pair of routers maintains a single TCP connection for all traffic that is sent between them. This includes multiplexed traffic for different circuits, and results in cross-circuit interference that degrades performance. We find that congestion control mechanisms are being unfairly applied to all circuits when they are intended to throttle only the noisy senders. We also show how packet dropping on one circuit causes interference on other circuits.

Section 3 presents our solution to this problem—a new transport layer that is backwards compatible with the existing Tor network. Routers in Tor can gradually and independently upgrade, and our system provides immediate benefit to any pair of routers that choose to use our improvements. It uses a separate TCP connection for each circuit, but secures the TCP header to avoid the disclosure of per-circuit data transfer

---

\*Work done while at the University of Waterloo

statistics. Moreover, it uses a user-level TCP implementation to address the issue of socket proliferation that prevents some operating systems from being able to volunteer as ORs.

Section 4 presents experiments to compare the existing Tor with our new implementation. We compare latency and throughput, and perform timing analysis of our changes to ensure that they do not incur non-negligible computational latency. Our results are favourable: the computational overhead remains negligible and our solution is successful in addressing the improper use of congestion control.

Section 5 compares our enhanced Tor to other anonymity systems, and Section 6 concludes with a description of future work.

## 1.1 Apparatus

Our experiments were performed on a commodity Thinkpad R60—1.66 GHz dual core with 1 GB of RAM. Care was taken during experimentation to ensure that the system was never under load significant enough to influence the results. Our experiments used a modified version of the Tor 0.2.0.x stable branch code.

## 2 Problems with Tor’s Transport Layer

We begin by briefly describing the important aspects of Tor’s current transport layer. For more details, see [3]. An end user of Tor runs an Onion Proxy on her machine, which presents a SOCKS proxy interface [10] to local applications, such as web browsers. When an application makes a TCP connection to the OP, the OP splits it into fixed-size *cells* which are encrypted and forwarded over a *circuit* composed of (usually 3) Onion Routers. The last OR creates a TCP connection to the intended destination host, and passes the data between the host and the circuit.

The circuit itself is constructed with hop-by-hop TCP connections, each protected with TLS [2], which provides confidentiality and data integrity. The OP picks a first OR ( $OR_1$ ), makes a TCP connection to it, and starts TLS on that connection. It then instructs  $OR_1$  to connect to a particular second OR,  $OR_2$ , of the OP’s choosing. If  $OR_1$  and  $OR_2$  are not already in contact, a TCP connection is established between them, again with TLS. If  $OR_1$  and  $OR_2$  are already in contact (because other users, for example, have chosen those ORs for their circuits), the existing TCP connection is used for all traffic between those ORs. The OP then instructs  $OR_2$  to contact a third OR,  $OR_3$ , and so on. Note that there is *not* an end-to-end TCP connection from the OP to the destination host, nor to any OR except  $OR_1$ .

This multi-hop transport obviously adds additional unavoidable latency. However, the observed latency of Tor is larger than accounted for simply by the additional transport time. In [15], the first author of this paper closely examined the sources of latency in a live Tor node. He found that processing time and input buffer queueing times were negligible, but that output buffer queueing times were significant. For example, 40% of output buffers had data waiting in them from 100 ms to over 1 s more than 20% of the time. The data was waiting in these buffers because the the operating system’s output buffer for the corresponding socket was itself full, and so the OS was reporting the socket as unwritable. This was due to TCP’s congestion control mechanism, which we discuss next.

Socket output buffers contain two kinds of data: packet data that has been sent over the network but is unacknowledged<sup>1</sup>, and packet data that has not been sent due to TCP’s congestion control. Figure 1 shows the size of the socket output buffer over time for a particular connection. First, unwritable sockets occur when the remaining capacity in an output buffer is too small to accept new data. This in turn occurs because there is already too much data in the buffer, which is because there is too much unacknowledged data in flight and throttled data waiting to be sent. The congestion window (CWND) is a variable that stores

---

<sup>1</sup>Recall that TCP achieves reliability by buffering all data locally until it has been acknowledged, and uses this to generate retransmission messages when necessary




Figure 1: TCP socket output buffer size, writability, and unacknowledged packets over time.

the number of packets that TCP is currently willing to send to the peer. When the number of packets in flight exceeds the congestion window then the sending of more data is throttled until acknowledgments are received. Once congestion throttles sending, the data queues up until either packets are acknowledged or the buffer is full.

In addition to congestion control, TCP also has a flow control mechanism. Receivers advertise the amount of data they are willing to accept; if more data arrives at the receiver before the receiving application has a chance to read from the OS's receive buffers, this advertised receiver window will shrink, and the sender will stop transmitting when it reaches zero. In none of our experiments did we ever observe Tor throttling its transmissions due to this mechanism; the advertised receiver window sizes never dropped to zero, or indeed below 50 KB. Congestion control, rather than flow control, was the reason for the throttling.

While data is delayed because of congestion control, it is foolhardy to attempt to circumvent congestion control as a means of improving Tor's latency. However, we observe that Tor's transport between ORs results in an *unfair* application of congestion control. In particular, Tor's circuits are multiplexed over TCP connections; i.e., a single TCP connection between two ORs is used for multiple circuits. When a circuit is built through a pair of unconnected routers, a new TCP connection is established. When a circuit is built through an already-connected pair of ORs, the existing TCP stream will carry both the existing circuits and the new circuit. This is true for all circuits built in either direction between the ORs.

In this section we explore how congestion control affects multiplexed circuits and how packet dropping and reordering can cause interference across circuits. We show that TCP does not behave optimally when circuits are multiplexed in this manner.

## 2.1 Unfair Congestion Control

We believe that multiplexing TCP streams over a single TCP connection is unwise and results in the unfair application of TCP's congestion control protocol. In particular, the nature of TCP's congestion control mechanism results in multiple data streams competing to send data over a TCP stream that gives more bandwidth to circuits that send more data (i.e., it gives each byte of data the same priority regardless of




Figure 2: Example of congestion on multiple streams.

its source). A busy circuit that triggers congestion control will cause low-bandwidth circuits to struggle to have their data sent. Figure 2 shows a demultiplexed buffer throughput graph for a real connection over Tor. Time increases along the X-axis, and data increases along the Y-axis. The two series correspond to two different streams that show when data is sent over time. The periods that swell indicate that Tor’s internal output buffer has swelled: the left edge grows when data enters the buffer, and the right edge grows when data leaves the buffer. This results in the appearance of a line when the buffer is well-functioning, and a triangular or parallelogram shape when data arrives too rapidly or the connection is troubled. Additionally, we strike a vertical line across the graph whenever a packet is dropped.

What we learn from this buffer is that it serves two circuits. One circuit serves one MB over ten minutes, and sends cells evenly. The other circuit is inactive for the most part, but three times over the execution it suddenly serves 200 KB of cells very quickly. We can see that each time this happens, the buffer swells with data, incurring a significant delay. Importantly, the other circuit is affected despite the fact that it did not change its behaviour. Congestion control mechanisms that throttle the TCP connection will give preference to the burst of writes because it simply provides more data, while the latency for a low-bandwidth application such as `ssh` increases unfairly.

## 2.2 Cross-Circuit Interference

Multiplexing multiple streams of data over a single TCP stream ensures that the received data will appear in the precise order in which the component streams were multiplexed. When packets are dropped or reordered, the TCP stack will buffer available data on input buffers until the missing in-order component is available. We hypothesize that when active circuits are multiplexed over a single TCP connection, Tor suffers an unreasonable performance reduction when either packet dropping or packet reordering occur. Cells may be available in-order for one particular circuit but are being delayed due to missing cells for another circuit. Ideally, multiplexed streams in TCP would be sent over the same connection using different channels, where in-order guarantees were only made for data sent over the same channel. However TCP sends all data over the same channel and so it is only readable in the order it was dispatched. Packet loss or reordering will




Figure 3: TCP correlated streams. Shades correspond to cells for different circuits.

Configuration	Network Throughput (KB/s)	Circuit Throughput (KB/s)	Throughput Degradation	Effective Drop Rate
No dropping	$221 \pm 6.6$	$36.9 \pm 1.1$	0 %	0 %
0.1 % (remaining)	$208 \pm 14$	$34.7 \pm 2.3$	6 %	0.08 %
0.1 % (shared)	$184 \pm 17$	$30.8 \pm 2.8$	17 %	0.03 %

Table 1: Throughput for different dropping configurations. Network throughput is the total data sent along all the circuits.

cause the socket to indicate that no data is available to read even if other circuits have their sequential cells available in buffers.

Figure 3 illustrates the classic head-of-line blocking behaviour of Tor during a packet drop; cells for distinct circuits are represented by shades and a missing packet is represented with a cross. We see that the white, light grey, and black circuits have had all of their data successfully received, yet the kernel will not pass that data to the Tor application until the dropped dark grey packet is retransmitted and successfully received.

We verify our cross-circuit interference hypothesis in two parts. First, in this section we show that packet drops on a shared link degrade throughput much more severely than drops over unshared links. Second, in Section 4, we will show that this effect disappears with our proposed solution. We performed Experiment 1 to investigate the effect of packet dropping on circuit multiplexing.

An important note regarding step 3: there is a difference between the frequency of traffic for the two sets of connections in our experiment (and all other experiments set up like Figure 4). Traffic for separate circuits that are multiplexed over a single connection can be sent together in a single TCP packet. Hence, the shared link may transport one packet that will scaffold multiple packets to send to the next router. Moreover, the remaining links span two hops along each circuit whereas the shared link spans only one. This means that each packet will be eligible to be dropped twice as often in the remaining links than the shared link. We elucidate this disparity by explicitly providing the effective drop rate; i.e., the ratio of packets dropped to the total number of packets.

The results of Experiment 1 are shown in Tables 1 and 2. The average results for throughput and delay were accumulated over half a dozen executions of the experiment, and the mean intervals for the variates

Configuration	Average Latency	Latency Increase	Effective Drop Rate
No dropping	$933 \pm 260$ ms	0 %	0 %
0.1 % (remaining)	$983 \pm 666$ ms	5.4 %	0.08 %
0.1 % (shared)	$1053 \pm 409$ ms	12.9 %	0.03 %

Table 2: Latency for different dropping configurations.

---

**Experiment 1** Determining the effect of packet dropping on circuit multiplexing.

---

- 1: A Tor network of six routers on a single host was configured to have a latency of 50 milliseconds and a variable packet drop rate.
  - 2: Eight clients built circuits that were fixed so that the second and third ORs were the same for each client, but the first hop was evenly distributed among the remaining routers. Figure 4 illustrates this setup.
  - 3: There were three runs of the experiment. The first did not drop any packets. The second dropped 0.1% of packets on the shared link, and the third dropped 0.1% of packets on the remaining links.
  - 4: The ORs were initialized and then the clients were run until circuits were established.
  - 5: Timing clients running locally connected to each OP and used Tor to tunnel their connection to a timing server running locally.
  - 6: Data was collected for one minute.
- 

are computed using Student’s T distribution to 95% confidence.

These results confirm our hypothesis. The throughput degrades nearly threefold when packets are dropped on the shared link instead of the remaining links. This is despite a significantly lower overall drop rate. The behaviour of one TCP connection can adversely affect all correlated circuits, even if those circuits are used to transport less data.

Table 2 suggests that latency increases when packet dropping occurs. Latency is measured by the time required for a single cell to travel alongside a congested circuit, and we average a few dozen such probes. Again we see that dropping on the shared link more adversely affects the observed delay despite a reduced drop rate. However, we note that the delay sees wide variance, and the 95% confidence intervals are quite large.

## 2.3 Summary

Multiplexing circuits over a single connection is a potential source of unnecessary latency since it causes TCP’s congestion control mechanism to operate unfairly towards connections with smaller demands on throughput. High-bandwidth streams that trigger congestion control result in low-bandwidth streams having their congestion window reduced unfairly. Packet dropping and reordering also cause available data for multiplexed circuits to wait needlessly in socket buffers. These effects degrade both latency and throughput, which we have shown in experiments.

To estimate the magnitude of this effect in the real Tor network, we note that 10% of tor routers supply 87% of the total network bandwidth [11]. A straightforward calculation shows that links between top routers—while only comprising 1% of the possible network links—transport over 75% of the data. At the time of writing, the number of OPs is estimated in the hundreds of thousands and there are only about one thousand active ORs [17]. Therefore, even while most users are idle, the most popular 1% of links will be frequently multiplexing circuits.

Ideally, we would open a separate TCP connection for every circuit, as this would be a more appropriate use of TCP between ORs; packet drops on one circuit, for example, would not hold up packets in other circuits. However, there is a problem with this naive approach. An adversary observing the network could easily distinguish packets for each TCP connection just by looking at the port numbers, which are exposed in the TCP headers. This would allow him to determine which packets were part of which circuits, affording him greater opportunity for traffic analysis. Our solution is to tunnel packets from multiple TCP streams over DTLS, a UDP protocol that provides for the confidentiality of the traffic it transports. By tunnelling TCP over a secure protocol, we can protect both the TCP payload and the TCP headers.

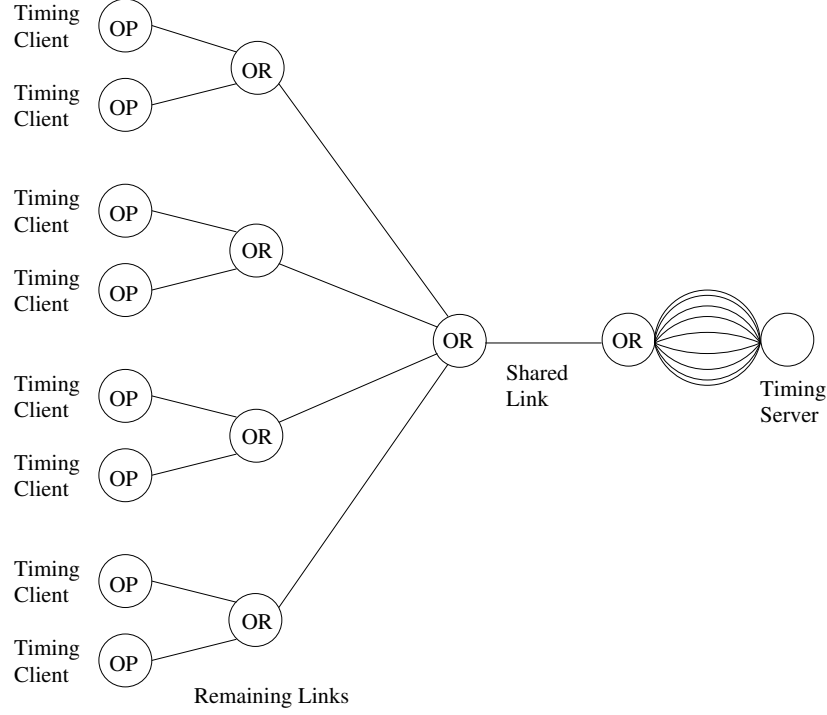


Figure 4: Setup for Experiment 1. The first two columns of links constitute the remaining links, in contrast to the shared link. The splay of links between the final OR and the timing server are each a separate TCP connection per circuit.

### 3 Proposed Transport Layer

This section proposes a TCP-over-DTLS tunnelling transport layer for Tor. Such a tunnel will transport TCP/IP packets between peers using DTLS—a secure datagram (UDP-based) transport [12]. A user-level TCP stack running inside Tor will be responsible for generating and parsing the TCP/IP packets that are sent over DTLS. Our solution will use a single unconnected UDP socket to communicate with all other ORs externally, and have a separate TCP connection for each circuit internally. This decorrelates circuits from TCP streams, which we have shown to be a source of unnecessary latency, while at the same time protects the TCP headers for separate streams from being visible to an attacker. Additionally, it reduces the number of sockets needed in kernel space, which is known to be a problem that prevents some Windows computers from volunteering as ORs. Figure 5 shows the design of our proposed transport layer, including how only a single circuit is affected by a dropped packet.

The interference that multiplexed circuits can have on each other during congestion, dropping, and reordering is a consequence of using a single TCP connection to transport data between each pair of ORs. This proposal uses a separate TCP connection for each circuit, ensuring that congestion or drops in one circuit will not affect other circuits. In order to hide the details of which traffic belongs to which circuit, and provide the necessary confidentiality and security of the transported cells, we use DTLS—secure datagram transport.

#### 3.1 A TCP-over-DTLS Tunnel

DTLS [12] is the datagram equivalent to the ubiquitous TLS protocol [2] that secures much traffic on the Internet today, including https web traffic, and indeed Tor. It provides for confidentiality and authenticity




Figure 5: Proposed TCP-over-DTLS Transport showing decorrelated streams. Shades correspond to cells for different circuits (cf. Figure 3).




Figure 6: Packets for TCP Tor and our TCP-over-DTLS improved Tor. Encrypted and authenticated components of the packet are shaded in grey. (a) shows the packet format for TCP Tor. (b) shows the packet format for our TCP-over-DTLS Tor. TORTP is a compressed form of the IP and TCP headers, and is discussed in Section 3.4.

of Internet datagrams, as well as other security properties such as replay prevention. IPsec [9] would have been another possible choice of protocol to use here; however, we chose DTLS for our application due to its acceptance as a standard and its existing implementation in the OpenSSL library (a library already in use by Tor). The TLS and DTLS APIs in OpenSSL are also unified; after setup, the same OpenSSL calls are used to send and receive data over either TLS or DTLS. This made supporting backwards compatibility easier: the Tor code will send packets either over TCP (with TLS) or UDP (with DTLS), as appropriate, with minimal changes.

Our new transport layer takes a prepared TCP/IP packet that was meant to be sent by the system, removes the unnecessary IP header, and encapsulates the TCP packet inside a DTLS packet that is then sent in a UDP/IP datagram. The receiving system will remove the UDP/IP header when receiving data from the socket, decrypt the DTLS payload to obtain a TCP packet, and translate it into a TCP/IP packet, which is then forwarded to the user-level TCP stack that processes the packet. A subsequent read from the user-level TCP stack will provide the packet data to our system.

In our system, the TCP sockets reside in user space, and the UDP sockets reside in kernel space. The use of TCP-over-DTLS affords us the great utility of TCP: guaranteed in-order delivery and congestion control. The user-level TCP stack provides the functionality of TCP, and the kernel-level UDP stack is used simply to transmit packets. The secured DTLS transport allows us to protect the TCP header from snooping and forgery and effect a reduced number of kernel-level sockets.

ORs require opening many sockets, and so our user-level TCP stack must be able to handle many concurrent sockets, instead of relying on the operating system’s TCP implementation which varies from system to system. In particular, some discount versions of Windows artificially limit the number of sockets the user can open, and so we use Linux’s free, open-source, and high-performance TCP implementation inside user space. Even Windows users will be able to benefit from an improved TCP implementation, and thus any user of an operating system supported by Tor will be able to volunteer their computer as an OR if they so choose.

UDP allows sockets to operate in an unconnected state. Each time a datagram is to be sent over the



Internet, the destination for the packet is also provided. Only one socket is needed to send data to every OR in the Tor network. Similarly, when data is read from the socket, the sender's address is also provided alongside the data. This allows a single socket to be used for reading from all ORs; all connections and circuits will be multiplexed over the same socket. When reading, the sender's address can be used to demultiplex the packet to determine the appropriate connection for which it is bound. What follows is that a single UDP socket can be used to communicate with as many ORs as necessary; the number of kernel-level sockets is constant for arbitrarily many ORs with which a connection may be established. This will become especially important as the number of nodes in the Tor network grows over time.

Figure 6(a) shows the packet format for TCP Tor, and Figure 6(b) shows the packet format for our TCP-over-DTLS Tor, which has expanded the encrypted payload to include the TCP/IP headers generated by the user-level TCP stack. The remainder of this section will discuss how we designed, implemented, and integrated these changes into Tor.

### 3.2 Backwards Compatibility

Our goal is to improve Tor to allow TCP communication using UDP in the transport layer. While the original ORs transported cells between themselves, our proposal is to transport, using UDP, both TCP headers and cells between ORs. The ORs will provide the TCP/IP packets to a TCP stack that will generate both the appropriate stream of cells to the Tor application, as well as TCP/IP packets containing TCP acknowledgements to be returned.

The integration of this transport layer into Tor has two main objectives. First is that of interoperability; it is essential that the improved Tor is backwards compatible with the TCP version of Tor so as to be easily accepted into the existing codebase. Recall that Tor has thousands of ORs, a client population estimated in the hundreds of thousands, and has not experienced any downtime since it launched in 2003. It is cumbersome to arrange a synchronized update of an unknown number of anonymous Tor users. A subset of nodes that upgrade and can take advantage of TCP-over-DTLS can provide evidence of the transport's improvement for the user experience—this incremental upgrade is our preferred path to acceptance. Our second objective is to minimize the changes required into the existing Tor codebase. We must add UDP connections into the existing datapath by reusing as much existing code as possible. This permits future developers to continue to improve Tor's datapath without having to consider two classes of communication. Moreover, it encourages the changes to quickly be incorporated into the main branch of the source code. While it will come with a performance cost for doing unnecessary operations, we perform timing analyses below to ensure that the resulting datapath latency remains negligible.

For interoperability we let older nodes communicate with the improved nodes using the same transport currently used. Newer nodes can advertise a UDP port for TCP-over-DTLS connections, but older nodes will continue to connect using the TCP port for OR communications. When the OR is instructed to make a connection (for example, when forming a circuit to a previously unconnected node) the connecting OR first checks if the destination OR offers a UDP service, using it if available. Thus, two UDP nodes will automatically communicate using UDP without disrupting the remaining nodes; their UDP connection will be inconspicuous even among other ORs involved in a circuit that uses that link. Traffic can switch between transport mechanisms in a node, since all traffic flows over a single connection between two nodes. All traffic that travels between a pair of UDP-capable nodes will then travel over a UDP link.

Clients of Tor will not need to upgrade their software to obtain the benefits of UDP transport. In fact, it is important that they continue to choose their circuit randomly among the ORs: intentionally choosing circuits consisting of UDP nodes when there are only a few such nodes decreases the privacy afforded to the client by rendering their circuit choices predictable.

### 3.3 User-level TCP Stack

If we simply replaced the TCP transport layer in Tor with a UDP transport layer, our inter-OR communication would then lack the critical features of TCP: guaranteed in-order transmission of streams, and the most well-studied congestion control mechanism ever devised. We wish to remove some of the unnecessary guarantees of TCP for the sake of latency; i.e., we do not need cells from separate circuits over the same connection to arrive in the order they were dispatched. However, we must still be able to reconstruct the streams of each individual circuit at both ends of the connection. We use a TCP implementation in user space (instead of inside the operating system) to accommodate us; a user-level TCP stack provides the implementation of the TCP protocols [6] as part of our program. User-level socket file descriptors and their associated data structures and buffers are accessible only in user-space and so are visible and relevant only to Tor. We use the UDP transport layer and DTLS to transport TCP/IP packets between the UDP peers. Only part of the TCP/IP packet is transmitted; the details will be discussed in section 3.4, but it serves our purposes now to conceptualize the two nodes as transporting full TCP/IP packets as the UDP datagram’s payload. Upon receiving a UDP datagram, the kernel will remove the UDP header and provide Tor with the enclosed DTLS packet; Tor will decrypt the DTLS payload and present the result (a TCP/IP packet) to its user-level TCP stack. Similarly, when the user-level TCP stack presents a packet for transmission, the node will encrypt it with DTLS and forward the resulting packet to the kernel which then sends it to the intended destination over UDP. The stack also performs retransmission and acknowledgement of TCP data that are integral to TCP’s reliability; these are protected with DTLS and forwarded over UDP in the same manner.

A user-level TCP stack provides an implementation of the suite of socket function calls, such as *socket()*, *send()*, and *recv()*. These reimplementations exist in harmony with the proper set of operating system commands, allowing both a user-level and kernel-level network layer. Thus, data structures and file descriptors created by calls to the user-level stack are visible and relevant only to the parent process; the operating system manages its sockets separately. The user-level stack responds to socket calls by generating packets internally for dispatching as dictated by TCP.

It may seem cumbersome to include an entire TCP implementation as a core component of Tor. In particular, patching the kernel’s implementation of TCP to support our features would take significantly less effort. However, Tor relies on volunteers to route traffic; complicated installation procedures are an immediate roadblock towards the ubiquitous use of Tor. The diverse operating systems Tor aims to support and the diverse skill level of its users prevent its installation from requiring external procedures, or even superuser privileges.

Daytona [14] is a user-level TCP stack that we chose for our purposes. It was created by researchers studying network analysis, and consists of the implementation of Linux’s TCP stack and the reimplementations of user-level socket functions. It uses `libpcap` to capture packets straight from the Ethernet device and a raw socket to write generated packets, including headers, onto the network. Daytona was designed to operate over actual networks while still giving user-level access to the network implementation. In particular, it allowed the researchers to tune the implementation while performing intrusive measurements. A caveat—there are licensing issues for Daytona’s use in Tor. As a result, the deployment of this transport layer into the real Tor network may use a different user-level TCP stack. Our design uses Daytona as a replaceable component and its selection as a user-level TCP stack was out of availability for our proof-of-concept.

### 3.4 UTCP: Our Tor-Daytona Interface

Our requirements for a user-level TCP stack are to create properly formatted packets, including TCP retransmissions, and to sort incoming TCP/IP packets into data streams: a black box that converts between streams and packets. For our purpose, all notions of routing, Ethernet devices, and interactions with a live network are unnecessary. To access the receiving and transmitting of packets, we commandeer the `rx()`

(receive) and `tx()` (transmit) methods of Daytona to instead interface directly with reading and writing to connections in Tor.

UTCP is an abstraction layer for the Daytona TCP stack used as an interface for the stack by Tor. Each UDP connection between ORs has a UTCP-connection object that maintains information needed by our stack, such as the set of circuits between those peers and the socket that listens for new connections. Each circuit has a UTCP-circuit object for similar purposes, such as the local and remote port numbers that we have assigned for this connection.

As mentioned earlier, only part of the TCP header is transmitted using Tor—we call this header the TORTP header. We do this simply to optimize network traffic, as much of the information in the 40 byte TCP/IP header is provided by the necessary UDP/IP header. The source and destination addresses and ports can be replaced with a numerical identifier that uniquely identifies the circuit for the connection. Since a UDP/IP header is transmitted over the actual network, Tor is capable of performing a connection look up based on the address of the packet sender. With the appropriate connection, and a circuit identifier, the interface to Daytona is capable of translating the TORTP header into the corresponding TCP/IP header.

When the UTCP interface receives a new packet, it uses local data and the TORTP headers to create the corresponding TCP/IP header. The resulting packet is then injected into the TCP stack. When Daytona’s TCP stack emits a new packet, a generic `tx()` method is invoked, passing only the packet and its length. We look up the corresponding UTCP circuit using the addresses and ports of the emitted TCP/IP header, and translate the TCP/IP header to our TORTP header and copy the TCP payload. This prepared TORTP packet is then sent to Tor, along with a reference to the appropriate circuit, and Tor sends the packet to the destination OR over the appropriate DTLS connection.

## 4 Experimental Results

In this section we perform experiments to compare the existing Tor transport layer with an implementation of our proposed TCP-over-DTLS transport. We begin by timing the new sections of code to ensure that we have not significantly increased the computational latency. Then we perform experiments on a local Tor network of routers, determining that our transport has indeed addressed the cross-circuit interference issues previously discussed.

### 4.1 Timing Analysis

Our UDP implementation expands the datapath of Tor by adding new methods for managing user-level TCP streams and UDP connections. We profile our modified Tor and perform static timing analysis to ensure that our new methods do not degrade the datapath unnecessarily. Experiment 2 was performed to profile our new version of Tor.

Our results indicate that no new datapath component results in a significant source of computational latency. We give our results to the eightieth percentile of measurements.

**Injecting** Injecting a new TCP packet is the part of the datapath that performs DTLS decryption, preprocessing, injecting the packet into the user-level TCP stack, and possibly invoking the transmit method to send an acknowledgement. The eightieth percentile measurement reports that injecting a new packet takes approximately 100 microseconds.

**Transmitting** Transmitting a new TCP packet is the part of the datapath that translates the packet header, performs DTLS encryption, and dispatches the result over the wire. The eightieth percentile measurement reports that transmitting a generated packet takes approximately 80 microseconds.

---

**Experiment 2** Timing analysis of our modified TCP-over-DTLS datapath.

---

- 1: TCP-over-DTLS Tor was modified to time the duration of the aspects of the datapath:
    - injection of a new packet,
    - emission of a new packet,
    - the TCP timer function, and
    - the entire datapath from reading a packet on a UDP socket, demultiplexing the result, injecting the packet, reading the stream, processing the cell, writing the result, and transmitting the generated packet.
  - 2: The local Tor network was configured to use 50 ms of latency between connections.
  - 3: A client connected through Tor to request a data stream.
  - 4: Data travelled through the network for several minutes.
- 

**TCP Timer** The TCP timer is not part of the main datapath, but is invoked every 10 milliseconds to perform some timer-based TCP operations: retransmission, delayed acknowledgements, etc. It increments the clock, checks for assigned work, and possibly invokes the transmission function. It only performs transmissions less than one percent of the time, so we only collect timing results for invocations where it actually performs work, which amortizes to an order of once per second. The eightieth percentile measurement for the TCP timer function is approximately 85 microseconds, which is because it performs many of the same operations as transmitting a packet.

**Datapath** The entire UDP datapath is the end-to-end operation: demultiplexing a UDP packet, injecting the data into the TCP stack (and possibly acknowledging the packet), reading from the user-level TCP stream, processing the read cells, writing the result to the socket, and possibly emitting a packet to the next hop. The eightieth percentile measurement is the time required to perform all these operations—the only additional operation that occurs at higher percentiles is the allocation of new sockets—and takes approximately 250 microseconds.

**Summary** The performance results indicate that our modifications to Tor are not going to contribute significantly to latency. The TCP timer is negligible. We have increased the datapath latency to an expected value of 250 microseconds per OR, or 1.5 milliseconds for a round trip along a circuit of length three. This is still an order of magnitude briefer than the round-trip times between ORs on a circuit (assuming geopolitically diverse circuit selection). Assuming each packet is 512 bytes (the size of a cell—a conservative estimate as our experiments have packets that carry full dataframes), we have an upper bound on throughput of 4000 cells per second or 2 MB/s. While this is a reasonable speed that will likely not form a bottleneck, Tor ORs that are willing to devote more than 2 MB/s of bandwidth may require better hardware than the Thinkpad R60 used in our experiments.

## 4.2 Basic Throughput

We perform Experiment 3 to compare the basic throughput and latency of our modification to Tor, the results of which are shown in Table 3. We can see that the UDP version of Tor has somewhat lower throughput. Negotiating the throughput up to this value took large amounts of TCP tuning and debugging the user-level TCP stack. In particular, errors were uncovered in Daytona’s congestion control implementation. While there may be slight degradation in performance when executing TCP operations in user-space instead of kernel-space, both implementations of TCP are based on the same Linux TCP implementation operating

---

**Experiment 3** Basic throughput and delay for TCP and TCP-over-DTLS versions of Tor.

---

- 1: A local Tor network running six routers on a local host was configured to have a latency of 50 milliseconds.
  - 2: The local network uses first the original TCP version of Tor, and then our modified TCP-over-DTLS version of Tor.
  - 3: Two OPs are configured to connect to our local Tor network along a fixed, shared circuit.
  - 4: The latency of the circuit is measured by sporadically requesting a timestamp along one circuit. The results indicate the base latency as there is no other data being sent through the network.
  - 5: The other circuit begins a large bulk transfer to measure throughput.
  - 6: Measurements for delay were collected in the presence of other traffic, and recorded separately from those of step 4.
  - 7: Data was collected for over a minute, and each configuration was run a half dozen times to obtain confidence intervals.
- 

Configuration	Network Throughput	Circuit Delay	Base Delay
TCP Tor	$176 \pm 24.9$ KB/s	$1026 \pm 418$ ms	$281 \pm 12$ ms
TCP-over-DTLS Tor	$111 \pm 10.4$ KB/s	$273 \pm 31$ ms	$260 \pm 1$ ms

Table 3: Throughput and delay for different reordering configurations. The base delay figures are measurements taken in the absence of other traffic, while the circuit delay figures are the measurements while bulk data transfer is occurring on another circuit.

over in the same network conditions. We would expect comparable throughputs as a result, and predict that with more effort to resolve outstanding bugs, or the integration of a user-level TCP stack better optimized for Tor’s needs, the disparity in throughputs will vanish. We discuss this further in the future work section.

More important is that the circuit delay for a second stream over the same circuit indicates that our UDP version of Tor vastly improves latency in the presence of a high-bandwidth circuit. When one stream triggers the congestion control mechanism, it does not cause the low-bandwidth client to suffer great latency as a consequence. In fact, the latency observed for TCP-over-DTLS is largely attributable to the base latency imposed on connections by our experimental setup. TCP Tor, in contrast, shows a three-and-a-half fold increase in latency when the circuit that it multiplexes with the bulk stream is burdened with traffic.

The disparity in latency for the TCP version means that information is leaked: the link between the last two nodes is witnessing bulk transfer. This can be used as a reconnaissance technique; an entry node, witnessing a bulk transfer from an client and knowing its next hop, can probe potential exit nodes with small data requests to learn congestion information. Tor rotates its circuits every ten minutes. Suppose the entry node notices a bulk transfer when it begins, and probes various ORs to determine the set of possible third ORs. It could further reduce this set by re-probing after nine minutes, after which time most of the confounding circuits would have rotated to new links.

We conclude that our TCP-over-DTLS, while currently suffering lower throughput, has successfully addressed the latency introduced by the improper use of the congestion control mechanism. We expect that once perfected, the user-level TCP stack will have nearly the same throughput as the equivalent TCP implementation in the kernel. The response latency for circuits in our improved Tor is nearly independent of throughput on existing Tor circuits travelling over the same connections; this improves Tor’s usability and decreases the ability for one circuit to leak information about another circuit using the same connection through interference.

Version	Configuration	Network Throughput (KB/s)	Circuit Throughput (KB/s)	Throughput Degradation	Effective Drop Rate
TCP-over-DTLS	No dropping	$284 \pm 35$	$47.3 \pm 5.8$	0 %	0 %
	0.1 % (remain.)	$261 \pm 42$	$43.5 \pm 7.0$	8 %	0.08 %
	0.1 % (shared)	$270 \pm 34$	$45.2 \pm 5.6$	4 %	0.03 %
TCP	No dropping	$221 \pm 6.6$	$36.9 \pm 1.1$	0 %	0 %
	0.1 % (remain.)	$208 \pm 14$	$34.7 \pm 2.3$	6 %	0.08 %
	0.1 % (shared)	$184 \pm 17$	$30.8 \pm 2.8$	17 %	0.03 %

Table 4: Throughput for different dropping configurations.

Version	Configuration	Average Latency	Latency Degradation	Effective Drop Rate
TCP-over-DTLS	No dropping	$428 \pm 221$ ms	0 %	0 %
	0.1 % (remaining)	$510 \pm 377$ ms	20 %	0.08 %
	0.1 % (shared)	$461 \pm 356$ ms	7 %	0.03 %
TCP	No dropping	$933 \pm 260$ ms	0 %	0 %
	0.1 % (remaining)	$983 \pm 666$ ms	5.4 %	0.08 %
	0.1 % (shared)	$1053 \pm 409$ ms	12.9 %	0.03 %

Table 5: Latency for different dropping configurations.

### 4.3 Multiplexed Circuits with Packet Dropping

Packet dropping occurs when a packet is lost while being routed through the Internet. Packet dropping, along with packet reordering, are consequences of the implementation of packet switching networks and are the prime reason for the invention of the TCP protocol. In this section, we perform an experiment to contrast the effect of packet dropping on the original version of Tor and our improved version.

We reperformed Experiment 1—using our TCP-over-DTLS implementation of Tor instead of the standard implementation—to investigate the effect of packet dropping. The results are presented in Tables 4 and 5. We reproduce our results from Tables 1 and 2 to contrast the old (TCP) and new (TCP-over-DTLS) transports.

We find that throughput is much superior for the TCP-over-DTLS version of Tor. This is likely because the TCP congestion control mechanism has less impact on throttling when each TCP stream is separated. One stream may back off, but the others will continue sending, which results in a greater throughput over the bottleneck connection. This is reasonable behaviour since TCP was designed for separate streams to function over the same route. If congestion is a serious problem then multiple streams will be forced to back off and find the appropriate congestion window. Importantly, the streams that send a small amount of data are much less likely to need to back off, so their small traffic will not have to compete unfairly for room inside a small congestion window intended to throttle a noisy connection. The benefits of this are clearly visible in the latency as well: cells can travel through the network considerably faster in the TCP-over-DTLS version of Tor. Despite the large confidence intervals for latency mentioned earlier, we see now that TCP-over-DTLS consistently has significantly lower latency than the original TCP Tor.

The TCP-over-DTLS version has its observed throughput and latency affected proportionally to packet drop rate. It did not matter if the drop was happening on the shared link or the remaining link, since the shared link is not a single TCP connection that multiplexes all traffic. Missing cells for different circuits no

longer cause unnecessary waiting, and so the only effect on latency and throughput is the effect of actually dropping cells along circuits.

## 5 Alternative Approaches

Tor is not the only anonymity system available. For comparison, we outline the mechanisms used by other anonymity systems to understand how they differ from Tor.

### 5.1 Anonymizer

Anonymizer [1] is a single-hop proxy. Clients open a connection to the anonymizing proxy, which strips their identifying information and forwards the result to the intended destination. This is a simple and effective design but the proxy is aware of the source and destination of messages—users must trust the single hop to protect their privacy. This creates a single point of failure; an entity that can observe traffic entering and exiting the proxy can break a user’s privacy. Since there are no circuits, each user has separate connections for all their traffic.

### 5.2 JonDonym

The JonDonym project [7], formerly known as the Java Anon Proxy, is a multi-hop proxy similar to both Tor and Anonymizer. It amalgamates traffic into cascades—a fixed well-known circuit. Multiple hops are used to ensure that no entity knows the source and destination of messages. However, the circuit is public; any observer knows the nodes it must compromise to break a user’s privacy. Users must trust the fact that the nodes on the fixed circuit do not collude. Moreover, an entity that can observe traffic at each end of the cascade can perform effective traffic analysis [3].

### 5.3 Freedom

The Freedom network [4], which was a business endeavour to provide Internet anonymity, shares many similarities to Tor. It uses client-guided circuit construction, where circuits are built through a network of routers whose operators were paid to route traffic. UDP was used as the transport layer for all traffic between nodes. TCP traffic was captured from the client and sent over UDP, forwarded along the circuit, and reassembled into TCP at the exit node. This means that the client and the exit node share a TCP connection, with end-to-end reliability and congestion control. This has been proposed for Tor as well, which is discussed next.

### 5.4 UDP-OR

Another similar transport mechanism for Tor has been proposed by Viecco [18] that encapsulates TCP packets from the OP and sends them over UDP until they reach the exit node. Reliability guarantees and congestion control are handled by the TCP stacks on the client and the exit nodes, and the middle nodes only forward traffic. A key difference between this approach and ours is that Viecco’s proposal requires a synchronized update of the Tor software for all users. This may be cumbersome given that Tor has thousands of routers and an unknown number of clients estimated in the hundreds of thousands.

This strategy proposes benefits in computational complexity and network behaviour. Computationally, the middle nodes must no longer perform unnecessary operations: packet injection, stream read, stream

write, packet generation, and packet emission. It also removes the responsibility of the middle node to handle retransmissions, which means a reduction in its memory requirements. The initial endpoint of communication will be responsible for retransmitting the message if necessary. We have shown that computational latency is insignificant in Tor, so this is simply an incidental benefit.

The tangible benefit of UDP-OR is to improve the network by allowing the ORs to function more exactly like routers. When cells arrive out of order at the middle node, they will be forwarded regardless, instead of waiting in input buffers until the missing cell arrives. Moreover, by having the sender's TCP stack view both hops as a single network, we alleviate problems introduced by disparity in network performance. Currently, congestion control mechanisms are applied along each hop, meaning that an OR in the middle of two connections with different performance metrics will need to buffer data to send over the slower connection. Tor provides its own congestion control mechanism, but it does not have the sophistication of TCP's congestion control.

We require experimentation to determine if this proposal is actually beneficial. While it is clear that memory requirements for middle nodes are reduced [18], the endpoints will see increased delay for acknowledgements. We expect an equilibrium for total system memory requirements since data will be buffered for a longer time. Worse, the approach shifts memory requirements from being evenly distributed to occurring only on exit nodes—and these nodes are already burdened with extra responsibilities. Since a significant fraction of Tor nodes volunteer only to forward traffic, it is reasonable to use their memory to ease the burden of exit nodes.

Circuits with long delays will also suffer reduced throughput, and so using congestion control on as short a path as possible will optimize performance. If a packet is dropped along the circuit, the endpoint must now generate the retransmission message, possibly duplicating previous routing efforts and wasting valuable volunteered bandwidth. It may be more efficient to have nodes along a circuit return their CWND for the next hop, and have each node use the minimum of their CWND and the next hop's CWND. Each node then optimizes their sending while throttling their receiving.

#### **5.4.1 Low-cost Privacy Attack**

UDP-OR may introduce an attack that permits a hostile entry node to determine the final node in a circuit. Previously each OR could only compute TCP metrics for ORs with whom they were directly communicating. Viecco's system would have the sender's TCP stack communicate indirectly with an anonymous OR. Connection attributes, such as congestion and delay, are now known for the longer connection between the first and last nodes in a circuit. The first node can determine the RTT for traffic to the final node. It can also reliably compute the RTT for its connection to the middle node. The difference in latency reflects the RTT between the second node and the anonymous final node. An adversary can use a simple technique to estimate the RTT between the second node and every other UDP Tor node in the network [5], possibly allowing them to eliminate many ORs from the final node's anonymity set. If it can reduce the set of possible final hops, other reconnaissance techniques can be applied, such as selectively flooding each OR outside of Tor and attempting to observe an increased latency inside Tor [13]. Other TCP metrics may be amalgamated to further aid this attack: congestion window, slow-start threshold, occurrence of congestion over time, standard deviation in round-trip times, etc. The feasibility of this attack should be examined before allowing nodes who do not already know each other's identities to share a TCP conversation.



## 6 Future Work

### 6.1 Real-World Benefits

We hope that our TCP-over-DTLS transport layer results in tangible improvements for the real Tor network, but we must begin to use it in the real world to test our hypothesis. After thorough testing, our transport layer can be introduced into the standard Tor release and Tor node operators can upgrade. Once a majority of users begin to use UDP connections, research can explore whether our proposal has been a benefit to the Tor network. Additionally, once most ORs use UDP, then we can determine if the reduced demand on open sockets solves the problem of socket proliferation on some operating systems.

### 6.2 TCP Stack Memory Management

Tor requires thousands of sockets to buffer fixed-size cells of data, but data is only buffered when it arrives out-of-order or has not been acknowledged. We envision dynamic memory management such as a shared cell pool to handle memory in Tor. Instead of repeatedly copying data cells from various buffers, each cell that enters Tor can be given a unique block of memory from the cell pool until it is no longer needed. A state indicates where this cell currently exists: input TCP buffer, input Tor buffer, in processing, output Tor buffer, output TCP buffer. This ensures that buffers are not allocated to store empty data, which reduces the overall memory requirements. Each cell also keeps track of its socket number, and its position in the linked list of cells for that socket. While each socket must still manage data such as its state and metrics for congestion control, this is insignificant as compared to the current memory requirements. This permits an arbitrary number of sockets, for all operating systems, and helps Tor’s scalability if the number of ORs increases by orders of magnitude.

This approach results in the memory requirements of Tor being a function of the number of cells it must manage at any time, independent of the number of open sockets. Since the memory requirements are inextricably tied to the throughput Tor offers, the user can parameterize memory requirements in Tor’s configuration just as they parameterize throughput. A client willing to denote more throughput than its associated memory requirements will have its contribution throttled as a result. If network conditions result in a surge of memory required for Tor, then it can simply stop reading from the UDP multiplexing socket. The TCP stacks that sent this unread data will assume there exists network congestion and consequently throttle their sending—precisely the behaviour we want—while minimizing leaked information about the size of our cell pool.

### 6.3 Stream Control Transmission Protocol

The Stream Control Transmission Protocol (SCTP) [16] is a message-based transport protocol. It provides similar features to TCP: connection-oriented reliable delivery with congestion control. However, it adds the ability to automatically delimit messages instead of requiring the receiving application to manage its own delimiters. The interface is based on sending and receiving messages rather than bytes, which is appropriate for Tor’s cell-based transport. As in our TCP-over-DTLS implementation, we would use SCTP-over-DTLS, securing the SCTP header and data inside a DTLS wrapper.

More importantly, SCTP also adds a feature well-suited to our purposes—multiple streams can be transported over the same connection. SCTP allows multiple independent ordered streams to be sent over the same socket; we can use this feature to send cells from different circuits on different streams. The in-order delivery guarantee is only provided for messages sent on the same stream, which is exactly the behaviour we want for cells from different circuits.

While SCTP is not as widely deployed as TCP, the concept of using a user-level SCTP implementation [8] suited for Tor remains feasible. Experimentation must be done to determine if SCTP’s congestion

control mechanism, which shares metrics and computations across all streams, acts fairly towards streams that send little data when compared to the streams that invoked congestion control.

## 7 Summary

Anonymous web browsing is an important step in the development of the Internet, particularly as it grows ever more inextricable from daily life. Tor is a privacy-enhancing technology that provides Internet anonymity using volunteers to relay traffic, and uses multiple relays in series to ensure that no entity (other than the client) in the system is aware of both the source and destination of messages.

Relaying messages increases latency since traffic must travel a longer distance before it is delivered. However, the observed latency of Tor is much larger than just this effect would suggest. To improve the usability of Tor, we examined where this latency occurs, and found that it happens when data sat idly in buffers due to congestion control. Since multiple Tor circuits are multiplexed over a single TCP connection between routers, we observed cross-circuit interference due to the nature of TCP's in-order, reliable delivery and its congestion control mechanisms.

Our solution was the design and implementation of a TCP-over-DTLS transport between ORs. Each circuit was given a unique TCP connection, but the TCP packets themselves were sent over the DTLS protocol, which provides confidentiality and security to the TCP header. The TCP implementation is provided in user-space, where it acts as a black box that translates between data streams and TCP/IP packets. We performed experiments on our implemented version using a local experimentation network and showed that we were successful in removing the observed cross-circuit interference and decreasing the observed latency.

## References

- [1] Anonymizer, Inc. Anonymizer FAQ. <http://www.anonymizer.com/company/about/anonymizer-faq.html>. Accessed February 2009.
- [2] Tim Dierks and Eric Rescorla. RFC 5246—The Transport Layer Security (TLS) Protocol Version 1.2. <http://www.ietf.org/rfc/rfc5246.txt>, August 2008.
- [3] Roger Dingledine, Nick Mathewson, and Paul Syverson. Tor: The Second-Generation Onion Router. *Proceedings of the 13th USENIX Security Symposium*, 2004.
- [4] Ian Goldberg and Adam Shostack. Freedom Network 1.0 Architecture and Protocols, 2001.
- [5] Krishna P. Gummadi, Stefan Saroiu, and Steven D. Gribble. King: Estimating Latency between Arbitrary Internet End Hosts. *ACM SIGCOMM Computer Communication Review*, 2002.
- [6] Information Sciences Institute. RFC 793—Transmission Control Protocol. <http://www.ietf.org/rfc/rfc793.txt>, September 1981.
- [7] JonDos GmbH. JonDonym—the internet anonymisation service. <https://www.jondos.de/en/>. Accessed February 2009.
- [8] Andreas Jungmaier, Herbert Hölzlzimmer, Michael Tüxen, and Thomas Dreibholz. The SCTP library (sctplib). <http://www.sctp.de/sctp-download.html>, 2007. Accessed February 2009.
- [9] Stephen Kent and Randall Atkinson. RFC 2401—Security Architecture for the Internet Protocol. <http://www.ietf.org/rfc/rfc2401.txt>, November 1998.

- [10] Marcus Leech et al. RFC 1928—SOCKS Protocol Version 5. <http://www.ietf.org/rfc/rfc1928.txt>, March 1996.
- [11] Damon McCoy, Kevin Bauer, Dirk Grunwald, Parisa Tabriz, and Douglas Sicker. Shining Light in Dark Places: A Study of Anonymous Network Usage. University of Colorado Technical Report CU-CS-1032-07, August 2007.
- [12] N Modadugu and E Rescorla. The Design and Implementation of Datagram TLS. *Network and Distributed System Security Symposium*, 2004.
- [13] Steven J. Murdoch and George Danezis. Low-Cost Traffic Analysis of Tor. In *IEEE Symposium on Security and Privacy*, pages 183–195, 2005.
- [14] P. Pradhan, S. Kandula, W. Xu, A. Shaikh, and E. Nahum. Daytona: A User-Level TCP Stack. <http://nms.lcs.mit.edu/~kandula/data/daytona.pdf>, 2002.
- [15] Joel Reardon. Improving Tor using a TCP-over-DTLS Tunnel. Master’s thesis, University of Waterloo, Waterloo, ON, September 2008.
- [16] Randall Stewart, Qiaobing Xie, Ken Morneault, Chip Sharp, Hanns Juergen Schwarzbauer, Tom Taylor, Ian Rytina, Malleswar Kalla, Lixia Zhang, and Vern Paxson. RFC 2960—Stream Control Transmission Protocol. <http://www.ietf.org/rfc/rfc2960.txt>, October 2000.
- [17] TorStatus. Tor Network Status. <http://torstatus.kgprog.com/>. Accessed February 2009.
- [18] Camilo Viecco. UDP-OR: A Fair Onion Transport Design. <http://www.petsymposium.org/2008/hotpets/udp-tor.pdf>, 2008.